TruCluster Server

Cluster Hardware Configuration Technical Update for the MSA1000

February 2004

Product Version: TruCluster Server Version 5.1A and 5.1B

Operating System and Version: Tru64 UNIX Version 5.1A and 5.1B

This technical update describes how to configure the StorageWorks Modular Smart Array 1000 (MSA1000) Fibre Channel storage system in a TruCluster Server environment.

Hewlett-Packard Company
Palo Alto, California

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Microsoft®, Windows®, and Windows NT® are trademarks of Microsoft Corporation in the U.S. and/or other countries. UNIX® and The Open Group™ are trademarks of The Open Group in the U.S. and/or other countries. All other product names mentioned herein may be trademarks of their respective owners.

Confidential computer software. Valid license from HP required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Contents

About This Technical Update

1 MSA1000 Fibre Channel Storage System

1.1 General Overview 1-1
1.2 Related Documents 1-1
1.3 MSA1000 Overview 1-2

2 Requirements and Restrictions

2.1 Supported AlphaServer Systems and Tru64 UNIX and TruCluster Server Software 2-1
2.2 General Requirements and Restrictions 2-2

3 TruCluster Server Configurations Using an MSA1000 for Storage

3.1 A Single Controller MSA1000 Cluster Configuration 3-1
3.2 No-Single-Point-of-Failure Cluster Configurations 3-3
3.3 Unsupported and Illegal Cluster Configurations 3-7

4 Preparing the MSA1000 for TruCluster Server Use

4.1 External Fibre Channel Switch Installation 4-2
4.2 Installing and Configuring the Fibre Channel Adapter Modules 4-2
4.2.1 Installing the Fibre Channel Adapter Modules 4-2
4.2.2 Setting the Fibre Channel Adapter to Run on a Fabric 4-3
4.2.3 Obtain the Fibre Channel Adapter Worldwide Name 4-6
4.3 Install and Set Up the MSA1000 Storage System 4-8
4.4 Installing Fiber-Optic Cables to the MSA1000 4-8
4.5 Installing Disk Enclosure Power Cables 4-9
4.6 Power on the MSA1000 Storage System 4-9
4.7 Prepare for CLI Use 4-9
4.8 Preparing the Storage for Tru64 UNIX and TruCluster Server Software Installation 4-10
4.9 Setting the Device Unit Number and Verifying Valid Boot Devices 4-22

4.9.1 Set the Device Unit Number 4-22
4.9.2 Displaying Valid Boot Devices 4-25
4.10 Installing the Software 4-27
4.10.1 Install the Base Operating System 4-27
4.10.2 Reset the bootdef_dev Console Environment Variable 4-28
4.10.3 Determining /dev/disk/dskn to Use for a Cluster Installation 4-30
4.10.4 Label the Disks to Be Used to Create the Cluster 4-32
4.10.5 Install the TruCluster Server Software and Create the First Cluster Member 4-32
4.10.6 Add an Additional System to the Cluster 4-32

A Worldwide Names, File System, Unit Number, Unit ID, Device Name, and Disk Name

Index

Examples

4-1 Verifying Fibre Channel Host Bus Adapter Topology 4-3
4-2 Correcting NVRAM Read Failed Message and Setting Fibre Channel Adapters to Run on Fabric 4-4
4-3 Changing the MSA1000 Controller Prompt 4-11
4-4 Displaying the MSA1000 Connections 4-11
4-5 Changing the Connection Profile to Tru64 4-13
4-6 Verifying the Change in Connection Name and Profile 4-13
4-7 Using the SHOW UNIT Command to Display Existing LUNs 4-14
4-8 Deleting a LUN 4-14
4-9 Displaying the Installed Disks 4-15
4-10 Creating the Storage Units (LUNs) 4-15
4-11 Adding a Spare Disk 4-17
4-12 Displaying Information About the Storage Unit 4-17
4-13 Changing a LUN Unit ID 4-20
4-14 Creating an Access Control List for TruCluster Server Storage Units 4-21
4-15 Verifying the Access Control List 4-21
4-16 Setting the Device Unit Number for the BOS Installation Disk 4-24
4-17 Setting the Device Unit Number for the First Cluster Member Boot Disk 4-24
4-18 Displaying the wwid and n Console Environment Variables ... 4-26

4-19 Sample Fibre Channel Device Names 4-26
4-20 Disks Available for Software Installation 4-27
4-21 Using the hwmgr to Display Devices with the Identifier Qualifier 4-31
4-22 Using the hwmgr to Display Devices with the MSA1000 Qualifier 4-31
4-23 Using the hwmgr to Display Devices with a Specific Unit ID . 4-31
4-24 Setting Device Unit Number for Additional Member System . 4-33

Figures

1-1 SCSI IDs for MSA1000 Disk Enclosure (Box 1) 1-3
1-2 SCSI Bus IDs for Model 4314R Disk Enclosures (Box 2 or Box 3) 1-3
1-3 SCSI Bus IDs for Model 4353R Disk Enclosure (Box 2) 1-4
3-1 Single-Point-of-Failure Configuration Using a Single Fibre Channel I/O Module 3-2
3-2 NSPOF TruCluster Server MSA1000 Configuration Using Fibre Channel Switches, Fibre Channel I/O Modules, and a Split-Bus Disk Enclosure 3-4
3-3 NSPOF TruCluster Server MSA1000 Configuration Using Fibre Channel Switches, Fibre Channel I/O Modules, and Single-Bus Disk Enclosures 3-5
3-4 NSPOF TruCluster Server MSA1000 Configuration Using Two MSA SAN Switch 2/8s 3-6
3-5 A TruCluster Server Configuration that is not Supported 3-7
3-6 An Illegal TruCluster Server Configuration 3-8

Tables

2-1 AlphaServer Systems, Minimum Software Versions, and SRM Firmware Required to Support the MSA1000 2-1
4-1 Fibre Channel Host Bus Adapter Port Worldwide Names 4-7
4-2 Storage Considerations for Example Installation 4-10
A-1 Fibre Channel Host Bus Adapter Worldwide Names A-1
A-2 Storage Allocation for TruCluster Server Installation A-1

About This Technical Update

This technical update provides important information about using the StorageWorks Modular Smart Array 1000 (MSA1000) Fibre Channel storage system in a TruCluster Server configuration.

Audience

If you plan to use an MSA1000 Fibre Channel storage system in a TruCluster configuration, read this addendum to the TruCluster Server *Cluster Hardware Configuration* manual.

Organization

This manual is organized as follows:

Chapter 1	Introduces the MSA1000 Fibre Channel storage system.
Chapter 2	Provides the requirements and restrictions to use an MSA1000 in a TruCluster Server configuration.
Chapter 3	Provides example supported, not supported, and illegal TruCluster Server configurations using an MSA1000 for storage.
Chapter 4	Describes how to prepare systems and storage for a TruCluster Server installation on MSA1000 storage system disks.
Appendix A	Provides blank tables to be used to log your Fibre Channel configuration.

1

MSA1000 Fibre Channel Storage System

This technical update to the TruCluster Server *Cluster Hardware Configuration* manual provides important information about support for the StorageWorks Modular Smart Array 1000 (MSA1000) Fibre Channel storage system designed for the entry-level to midrange storage area network (SAN). It has been qualified with the Tru64 UNIX and TruCluster Server Version 5.1A and 5.1B products in a TruCluster Server configuration.

Note

The name for the MSA1000 has been chanaged from Modular SAN Array to Modular Smart Array. The MSA1000 maintenance and service guides still use the old name.

1.1 General Overview

The MSA1000 is qualified for use with the TruCluster Server Version 5.1A and 5.1B products.

There is no TruCluster Server software release that coincides with the availability of support for the MSA1000, so this technical update provides the configuration information that would appear in the TruCluster Server *Cluster Hardware Configuration* manual.

1.2 Related Documents

For more information on the MSA1000 Fibre Channel storage system, see the following StorageWorks documentation:

- *hp StorageWorks Modular Smart Array 1000 Specific for Tru64 UNIX or OpenVMS Only QuickSpecs*
- *hp StorageWorks Modular Smart Array 1000 (MSA1000) Release Notes*
- *hp StorageWorks modular SAN array 1000 user guide*
- *hp StorageWorks modular SAN array 1000 maintenance and service guide*

1.3 MSA1000 Overview

The MSA1000 Fibre Channel storage array provides 2-Gb connections to the host with support for 1-Gb Fibre Channel connections. The MSA1000 consists of a 4U (17.6 cm (7.0 inches)) rack mount controller shelf with an attached disk enclosure that holds up to 14 1-inch Ultra3 SCSI disk drives.

The MSA1000 provides the following RAID support:

- RAID0: Striping, no fault tolerance. See RAID1+0.
- RAID1: Mirroring. See RAID1+0.
- RAID1+0: Striping and Mirroring. When a multiple of two drives are mirrored, data is automatically stripped within each mirrorset.
- RAID5: Distributed Data Guarding
- RAID ADG: Advanced Data Guarding

MSA1000 part numbers may be found in the *hp StorageWorks Modular Smart Array 1000 Specific for Tru64 UNIX or OpenVMS Only QuickSpecs*.

The total amount of disk space can be expanded by the addition of up to two Model 4314R or one Model 4354R StorageWorks disk enclosures (each of which holds 14 disks). An MSA1000 Fibre Channel storage system with two Model 4314R disk enclosures provides a maximum of 42 disk drives. An MSA1000 storage system with an optional Model 4354R disk enclosure provides a maximum of 28 disks.

The controller shelf provides space for 1 or 2 MSA1000 controllers, each with its associated input medium. The input medium may be:

- An embedded single input 2-Gb Fibre Channel I/O module that is part of a Fibre Channel switch fabric. The Fibre Channel I/O module is connected to a Fibre Channel switch, and it provides an input to the MSA1000 controller through the direct connection.
- An embedded 2-Gb MSA SAN Switch 2/8. This 8-port Fibre Channel switch has 7 external ports and 1 internal port. The internal port is connected directly to the MSA1000 controller.

Note

If there are two MSA1000 controllers, there must be two input modules of the same type, one for each MSA1000 controller.

The MSA1000 is compatible with 1-Gb or 2-Gb Fibre Channel fabrics.

The MSA1000 provides Ultra3 SCSI connections to the Model 4314R or 4354R disk enclosures, which house universal 1-inch Ultra3 hard drives. There may be a maximum of four SCSI buses: two internal SCSI buses for the MSA1000 enclosure (SCSI buses 0 and 1), and two external SCSI buses in the disk enclosures served by the SCSI I/O module (SCSI buses 2 and 3).

The MSA1000 drive enclosure has a split bus with two SCSI buses: SCSI bus 0 serves bays 1 through 7, and SCSI bus 1 serves bays 8 through 14.

The SCSI I/O module provides SCSI bus 2 and 3, which can be connected to optional expansion Model 4314R (a maximum of two) or Model 4354R (a maximum of one) StorageWorks disk enclosures.

The correlation between bay number, SCSI bus number, and SCSI bus ID for the MSA1000 disk enclosure, Model 4314R, and Model 4354R disk enclosures are shown in Figure 1-1, Figure 1-2, and Figure 1-3.

Figure 1-1: SCSI IDs for MSA1000 Disk Enclosure (Box 1)

BAY	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI Bus	SCSI Bus 0							SCSI Bus 1						
SCSI ID	0	1	2	3	4	5	8	0	1	2	3	4	5	8

ZK-2007U-AI

Figure 1-2: SCSI Bus IDs for Model 4314R Disk Enclosures (Box 2 or Box 3)

BAY	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI Bus	SCSI Bus 2 or SCSI Bus 3													
SCSI ID	0	1	2	3	4	5	8	9	10	11	12	13	14	15

ZK-2008U-AI

Figure 1-3: SCSI Bus IDs for Model 4353R Disk Enclosure (Box 2)

BAY	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI Bus	SCSI Bus 2							SCSI Bus 3						
SCSI ID	0	1	2	3	4	5	8	0	1	2	3	4	5	8

ZK-2009U-AI

2

Requirements and Restrictions

This chapter describes the hardware requirements and restrictions for use of an MSA1000 Fibre Channel storage array in a TruCluster Server cluster. It includes lists of supported Tru64 UNIX operating system and TruCluster Server versions, AlphaServer systems, standard reference manual (SRM) firmware, Fibre Channel host bus adapters, and other general restrictions.

2.1 Supported AlphaServer Systems and Tru64 UNIX and TruCluster Server Software

The MSA1000 is supported on the AlphaServer systems, with the minimum versions of Tru64 UNIX and TruCluster Server software and system reference manual (SRM) firmware shown in Table 2-1.

Table 2-1: AlphaServer Systems, Minimum Software Versions, and SRM Firmware Required to Support the MSA1000

AlphaServer Systems	Minimum Version of Required Software	SRM Firmware Version
DS10, DS10L, DS20E, DS25, ES40, ES45, GS80, GS160, and GS320	Version 5.1A with PK4 and NHD6 (or later)[a]	6.4
ES47, ES80, GS1280	Version 5.1B with PK1 (or later)[a]	6.5-8

[a] NHD7 includes Version 5.1A PK6 and Version 5.1B PK3, both of which are cumulative patches.

NHD7 is available for download at the following URL:

http://www2.itrc.hp.com/service/patch/mainPage.do

After logging in, you will be on the patch database home Web page. Locate, and download NHD7 as follows:

- Enter NHD7 in the find a specific patch window and click on the arrows (») to the right of the window.
- Click on any NHD7 in the table.
- Scroll down the page and click on add to selected patch list ».
- On the selected patch list page, click on download selected ».

- On the download patches page:
  - Select a server in your region.
  - Select the desired format for the download (zip, gzip, tar, or download a script that will ftp the patches).
  - Click on the download » button.

The MSA1000 is supported in a TruCluster Server cluster by AlphaServer systems using the following Fibre Channel host bus adapters:

- DS-KGPSA-CA (1-Gb) — DS10, DS10L, DS20E, DS25, ES40, ES45, GS80, GS160, and GS320
- DS-KGPSA-DA (2-Gb, FCA2354) — DS10, DS10L, DS20E, DS25, ES40, ES45, ES47, ES80, GS80, GS160, GS320, and GS1280
- DS-KGPSA-EA (2-Gb, FCA2384) — DS10, DS10L, DS20E, DS25, ES40, ES45, ES47, ES80, GS80, GS160, GS320, and GS1280

2.2 General Requirements and Restrictions

This section covers the requirements and restrictions for use of the MSA1000 in a TruCluster Server configuration.

For additional restrictions, or changes, and firmware upgrade procedures, refer to the *hp StorageWorks Modular Smart Array 1000 (MSA1000) Release Notes* at the MSA1000 home page at the following URL:

http://www.hp.com/go/msa1000/

- Select software, firmware, & drivers in the product information list.
- Select MSA1000 Alternative Firmware Tru64 UNIX, OpenVMS, or Alpaha-based servers under Controller Firmware.
- Select Alpha MSA1000/Tru64 UNIX/OpenVMS Firmware v4.24 - Zip file format .
- Extract MSA1000ReleaseNotes.pdf, the MSA1000 release notes.

If there are differences between this technical update and the MSA1000 release notes, the release notes have precedence over this technical update.

The requirements and restrictions for use of the MSA1000 in a TruCluster Server configuration are as follows:

- At the time of publication, the MSA1000 is being shipped with the latest firmware, version 4.24. If your MSA1000 does not have the latest firmware, it must be upgraded. Refer to the MSA1000 release notes for firmware upgrade instructions.

- TruCluster Server supports clusters with up to eight member systems. An MSA1000 may be connected to no more than four of these cluster member systems.
- Fibre Channel arbitrated loop (FC-AL), including the MSA Hub 2/3, is not supported with the MSA1000 at this time.
- The MSA Fabric Switch 6 (with five external and one internal ports) is not supported at this time.
- If you install a second controller, you must also install a second Fibre Channel I/O module or MSA SAN Switch 2/8.
- MSA1000 controller failover is not fully supported at this time. Only fabric failover is fully supported. Refer to the *hp StorageWorks Modular Smart Array 1000 (MSA1000) Release Notes* for more information.
- An MSA1000 with dual-redundant controllers operates in Active/Passive mode; one controller is active while the other controller is passive.
- The command-line interface (CLI) is used to configure and manage the MSA1000.
- The maximum length of the optical cable between the Fibre Channel host bus adapter and the switch (or hub), or the switch (or hub) and the MSA1000, is 500 meters (1640.4 feet) via shortwave multimode Fibre Channel cable at 1-Gb/sec. The distance is shortened to 300 meters (924 feet) at 2-Gb/sec.
- Only the Model 4314R (DS-SL13R-AA, 190209-001) and Model 4354R (DS-SL13R-BA, 190211-001) storage enclosures are supported.
- Two Model 4314R storage enclosures may be connected to an MSA1000.
- Only one Model 4354R storage enclosure may be connected to an MSA1000.
- The following 68-pin very high density cable interconnect (VHDCI-to-VHDCI) cables are supported between the MSA1000 SCSI I/O module and the Model 4314R or 4354R disk enclosure I/O modules:
  - 3X-BC56J-02, 341174-B21 (1.8 m (6 feet))
  - 3X-BC56J-03, 341175-B21 (3.7 meters (12 feet))
  - 3X-BC56J-04, 164604-B21 (7.3 meters (24 feet))

3

TruCluster Server Configurations Using an MSA1000 for Storage

This chapter provides illustrations of TruCluster Server configurations using an MSA1000 Fibre Channel storage array.

3.1 A Single Controller MSA1000 Cluster Configuration

A TruCluster Server configuration with a single-controller MSA1000 is shown in Figure 3-1. In this single controller configuration, the single Fibre Channel host bus adapter, Fibre Channel switch, MSA1000 controller, and Fibre Channel I/O module are each single points of failure.

Figure 3-1: Single-Point-of-Failure Configuration Using a Single Fibre Channel I/O Module

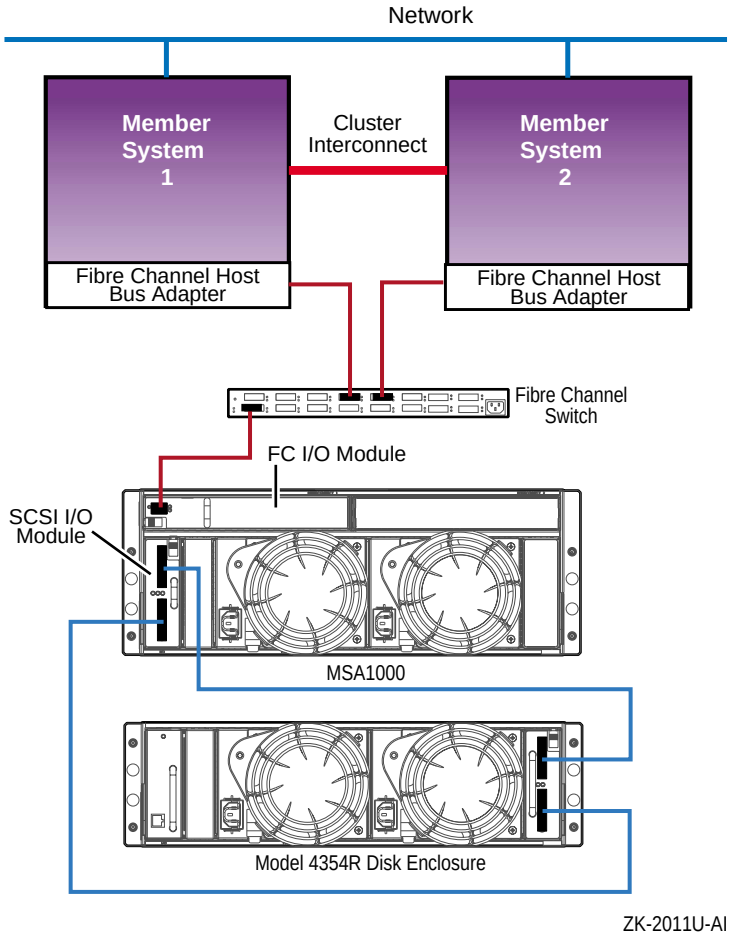

If a second MSA1000 controller and associated Fibre Channel I/O module is installed in the configuration shown in Figure 3-1, the Fibre Channel switch and single Fibre Channel host bus adapter in each system is still a single-point-of-failure.

3.2 No-Single-Point-of-Failure Cluster Configurations

Figure 3-2 is a no-single-point-of-failure (NSPOF) TruCluster Server configuration using two MSA1000 controllers. Each MSA1000 controller has a corresponding Fibre Channel I/O module. This configuration uses two Fibre Channel switches in a dual fabric to provide the input to the Fibre Channel I/O modules. A Model 4354R disk enclosure provide extra disk storage.

Note

Tru64 UNIX cannot detect all MSA1000 controller failover conditions. In these cases, Tru64 UNIX is incapable of forcing the other controller to take over. Therefore, an NSPOF configuration is not a true NSPOF configuration under all conditions. Refer to the *hp StorageWorks Modular Smart Array 1000 (MSA1000) Release Notes* for more information.

Figure 3-2: NSPOF TruCluster Server MSA1000 Configuration Using Fibre Channel Switches, Fibre Channel I/O Modules, and a Split-Bus Disk Enclosure

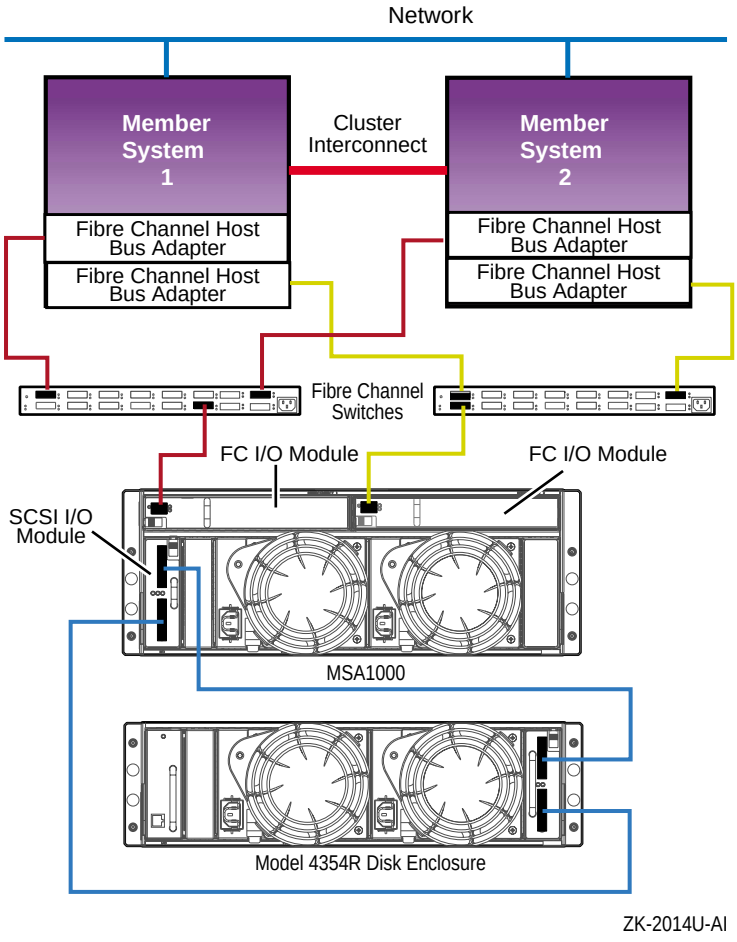

Figure 3-3 is the same as Figure 3-2, except for the use of two single-bus Model 4314R disk enclosures for additional storage.

Figure 3-3: NSPOF TruCluster Server MSA1000 Configuration Using Fibre Channel Switches, Fibre Channel I/O Modules, and Single-Bus Disk Enclosures

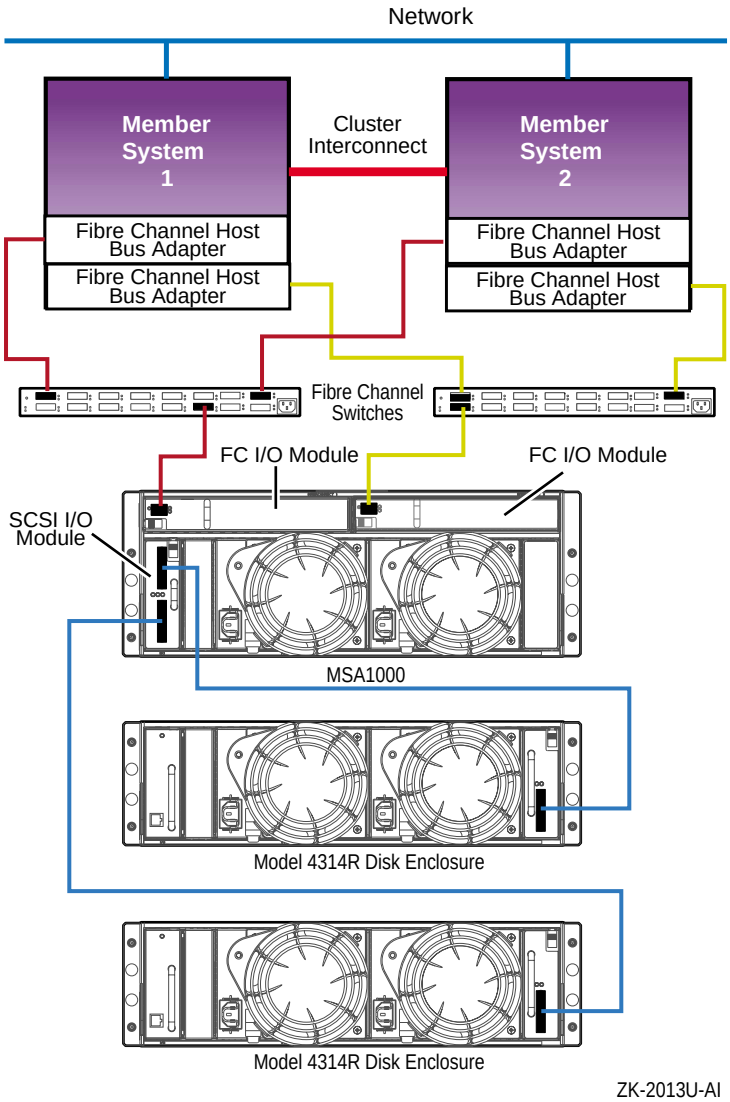

Figure 3-4 is an NSPOF TruCluster Server configuration using an MSA1000 with two controllers and MSA SAN Switch 2/8 Fibre Channel switches as inputs to the controllers. An MSA SAN Switch 2/8 is required for each controller.

Figure 3-4: NSPOF TruCluster Server MSA1000 Configuration Using Two MSA SAN Switch 2/8s

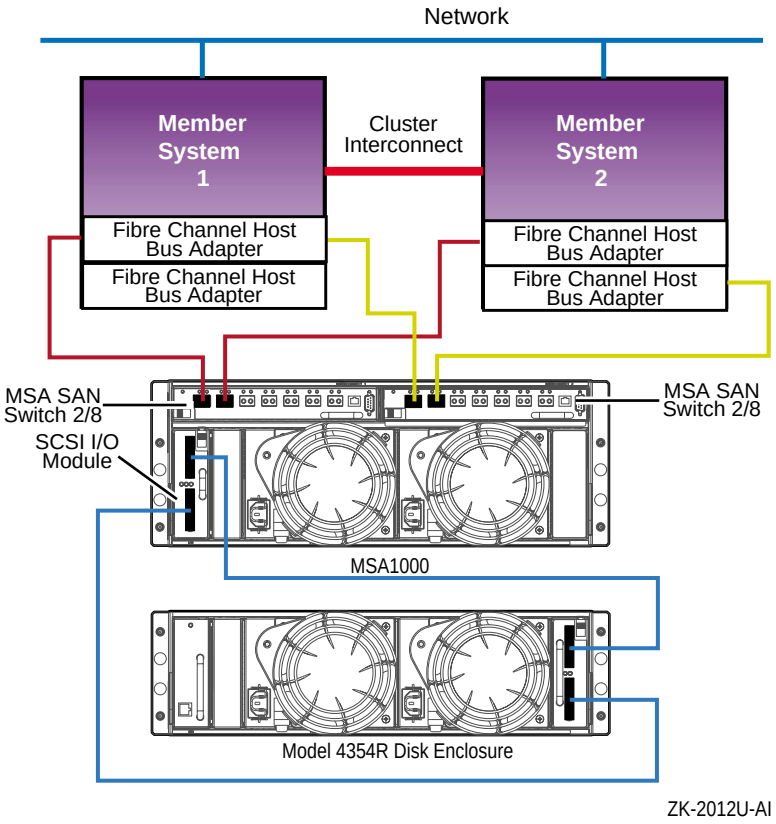

3.3 Unsupported and Illegal Cluster Configurations

Figure 3-5 shows a TruCluster Server configuration that is not supported. The MSA1000 controllers operate in active/passive mode; one controller is active while the other controller is passive. In the configuration shown in Figure 3-5, both cluster member systems try to make their controller active, creating a ping-pong effect. Both systems would receive many errors, and no useful I/O would be accomplished.

Figure 3-5: A TruCluster Server Configuration that is not Supported

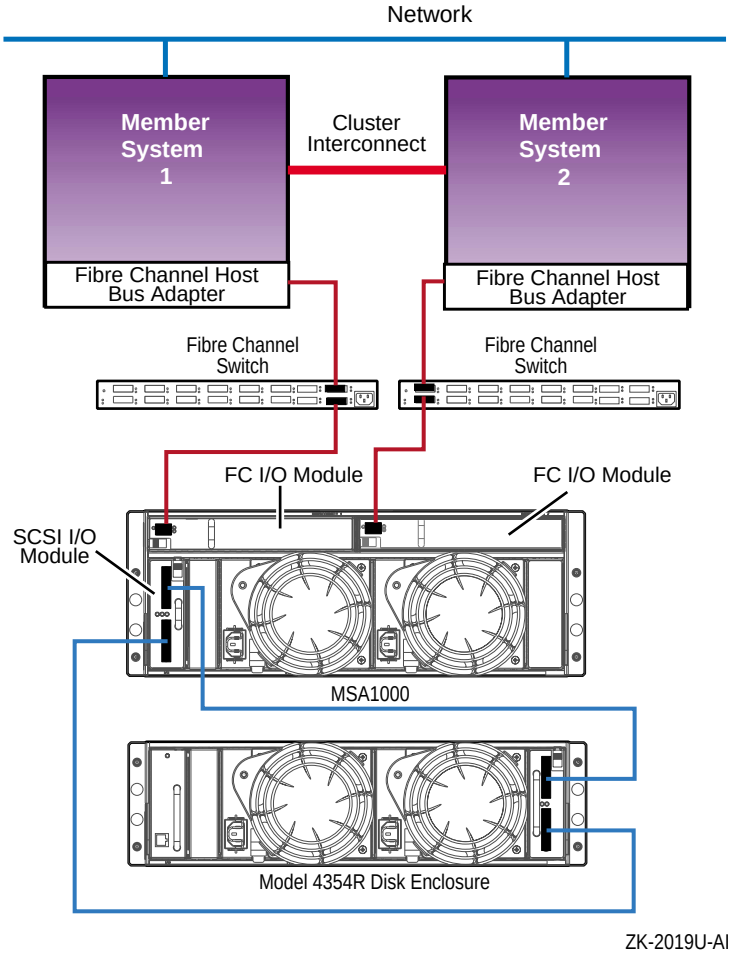

The configuration shown in Figure 3-6 is an illegal configuration. Tru64 UNIX requires that the Fibre Channel host bus adapter be connected to a Fibre Channel switch; the Fibre Channel host bus adapter cannot be connected directly to the array controller or Fibre Channel I/O module.

Figure 3-6: An Illegal TruCluster Server Configuration

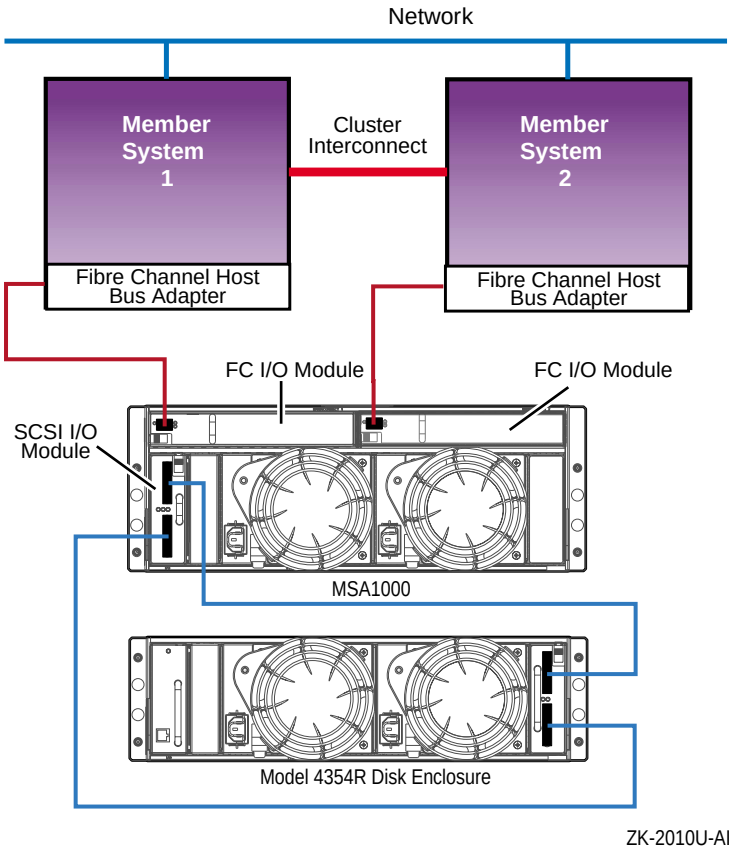

4

Preparing the MSA1000 for TruCluster Server Use

The topics in this chapter provide information on how to prepare an MSA1000 Fibre Channel storage system and associated disks for Tru64 UNIX and TruCluster Server installation.

The base operating system and TruCluster Server software will be installed on MSA1000 storage.

To install Tru64 UNIX and TruCluster Server on MSA1000 disks, follow these steps:

1. If the MSA1000 has installed Fibre Channel I/O modules, install the external Fibre Channel switches (Section 4.1).

   If the MSA1000 uses the embedded MSA1000 SAN switch 2/8, instead of the FC I/O module, ensure that they are installed
2. Install and configure the Fibre Channel host bus adapters (Section 4.2).
3. Install and set up the MSA1000 (Section 4.3).
4. Install fiber-optic cables (Section 4.4).
5. Install storage shelf power cables (Section 4.5).
6. Power on the MSA1000 storage system (Section 4.6).
7. Prepare a laptop computer or terminal for command-line interface use (Section 4.7).
8. If necessary, upgrade the MSA1000 firmware. The firmware is available from the MSA1000 home page at:

   `http://www.hp.com/go/msa1000/`

   See Section 2.2 for more information.
9. Set up storage for Tru64 UNIX and TruCluster Server installation (Section 4.8).
10. Set the device unit number and verify valid boot devices (Section 4.9).
11. Install Tru64 UNIX and TruCluster Server software (Section 4.10).

4.1 External Fibre Channel Switch Installation

If the MSA1000 has an installed Fibre Channel I/O module, you must use external Fibre Channel switches. An NSPOF configuration requires redundant MSA1000 controllers and two Fibre Channel I/O modules.

Ensure that the member systems, the Fibre Channel switches, and the MSA1000 are placed within the lengths of the optical cables that you will be using.

Install and set up your Fibre Channel switches using the documentation that came with the switch as follows:

- Install a minimum of two Fibre Channel switches if you have plans for a no-single-point-of-failure (NSPOF) configuration.
- All switches have a 10Base-T Ethernet (RJ45) port, and after the IP address is set, the Ethernet connection allows you to manage the switch:
  - Remotely using a telnet TCP/IP connection
  - With the Simple Network Management Protocol (SNMP)
  - Using Web management tools

If it is necessary to set up switch zoning, you can do so after installing the Fibre Channel host bus adapters, storage hardware, and associated cabling.

4.2 Installing and Configuring the Fibre Channel Adapter Modules

The following sections discuss Fibre Channel adapter (FCA) installation and configuration.

4.2.1 Installing the Fibre Channel Adapter Modules

To install the DS-KGPSA-CA, DS-KGPSA-DA (FCA2354), or DS-KGPSA-EA (FCA2384) Fibre Channel host bus adapter modules, follow these steps. For more information, refer to the following documentation:

- *64-Bit PCI-to-Fibre Channel Host Bus Adapter User Guide*
- *Tru64 UNIX and OpenVMS FCA-2354 Host Bus Adapter Installation Guide*
- *hp StorageWorks fca2384 PCI-X host bus adapter for OpenVMS and Tru64 UNIX installation guide*

Caution

Static electricity can damage modules and electronic components. We recommend using a grounded antistatic wrist strap and a grounded work surface when handling modules.

1. Using the documentation that came with your Fibre Channel host bus adapter, install the Fibre Channel adapter in an open 64-bit PCI slot.
2. Set the Fibre Channel adapter to run on fabric (see Section 4.2.2)
3. Obtain the Fibre Channel adapter worldwide node and worldwide port names (see Section 4.2.3).

4.2.2 Setting the Fibre Channel Adapter to Run on a Fabric

The Fibre Channel host bus adapter (FCA) defaults to the fabric mode, and can be used in a fabric without taking any action. However, if you install an FCA that has been used in the loop mode on another system, you will need to reformat the nonvolatile RAM (NVRAM) and configure it to run in a Fibre Channel fabric configuration.

Use the worldwide ID manager (`wwidmgr`) console utility to determine the mode of operation of the Fibre Channel host bus adapter, and to set the mode if it needs changing (for example, from loop to fabric).

Use the `wwidmgr` utility to verify that the topology for all Fibre Channel host bus adapters are set to fabric as shown in Example 4-1 and Example 4-2.

Example 4-1: Verifying Fibre Channel Host Bus Adapter Topology

```
P00>>> wwidmgr -show adapter
Link is down.
item       adapter           WWN                Cur. Topo  Next Topo
pga0.0.0.3.1 - Nvram read failed
[ 0] pga0.0.0.2.0         1000-0000-c922-4aac    FABRIC     UNAVAIL
pgb0.0.0.4.0 - Nvram read failed
[ 1] pgb0.0.0.4.0         1000-0000-c924-4b7b    FABRIC     UNAVAIL
[9999] All of the above.
```

A `Link is down` message indicates that one of the adapters is not available, which means that its probably not plugged into a switch.

The warning message `Nvram read failed` indicates that the Fibre Channel host bus adapter nonvolatile random-access memory (NVRAM) has not been initialized and formatted. The next topology will always be `UNAVAIL` for the host bus adapter that has an unformatted NVRAM. Both messages are benign and can be ignored for the fabric mode of operation.

The display in Example 4-1 shows that both Fibre Channel host bus adapters are set for fabric topology as the current topology, the default. When operating in a fabric, if the current topology is FABRIC, it does not matter if the next topology is Unavail, or that the NVRAM is not formatted (Nvram read failed).

You may view the Nvram read failed message as an error, and may not like the fact that the next topology is Unavail.

To correct the Nvram read failed situation and set the next topology to fabric, use the wwidmgr -set adapter command as shown in Example 4-2. This command initializes the NVRAM and sets the mode of all Fibre Channel host bus adapters to fabric.

Example 4-2: Correcting NVRAM Read Failed Message and Setting Fibre Channel Adapters to Run on Fabric

```
P00>>> wwidmgr -set adapter -item 9999 -topo fabric
Reformatting nvram
Reformatting nvram
P00>>> init
```

Note

The qualifier in the previous command is -topo and not -topology. You will get an error if you use -topology.

The item 9999 qualifier sets all Fibre Channel adapters. To set a single Fibre Channel adapter, use the item qualifier with that Fibre Channel adapter's item number (see Example 4-1).

If, for some reason, the current topology is LOOP, you have to change the topology to FABRIC to operate in a fabric. You will never see the Nvram read failed message if the current topology is LOOP. The NVRAM has to have been formatted to change the current mode to LOOP.

Consider the case where the Fibre Channel host bus adapter current topology is LOOP as follows:

```
P00>>> wwidmgr -show adapter
item       adapter           WWN                Cur. Topo  Next Topo
[ 0] pga0.0.0.2.0         1000-0000-c922-4aac    LOOP       LOOP
[ 1] pgb0.0.0.4.0         1000-0000-c924-4b7b    LOOP       LOOP
[9999] All of the above.
```

If the current topology for an adapter is LOOP, set an individual adapter to FABRIC by using the item number for that adapter (for example, 0 or 1). Use 9999 to set all adapters as shown in Example 4-2.

Displaying the adapter information again will show the topology that the adapters will assume after the next console initialization:

```
P00>>> wwidmgr -show adapter
item    adapter                  WWN           Cur. Topo  Next Topo
[ 0] pga0.0.0.2.0         1000-0000-c922-4aac    LOOP       FABRIC
[ 1] pgb0.0.0.4.0         1000-0000-c924-4b7b    LOOP       FABRIC
[9999] All of the above.
```

This display shows that the current topology for both Fibre Channel host bus adapters is LOOP, but will be FABRIC after the next initialization:

```
P00>>> init
P00>>>  wwidmgr -show adapter
item    adapter                  WWN           Cur. Topo  Next Topo
[ 0] pga0.0.0.2.0         1000-0000-c922-4aac    FABRIC     FABRIC
[ 1] pgb0.0.0.4.0         1000-0000-c924-4b7b    FABRIC     FABRIC
[9999] All of the above.
```

Notes

The console remains in wwid manager mode, and you cannot boot until the system is reinitialized. Use the init command or a system reset to reinitialize the system after using the wwid manager.

If you try to boot the system and receive the following error, initialize the console to get out of WWID manager mode and reboot:

```
P00>>> boot
warning -- main memory zone is not free
P00>>> init
       .
       .
       .
P00>>> boot
```

If you shut down the operating system and try to use the wwidmgr utility, you may be prevented from doing so. If you receive the following error, initialize the system and retry the wwidmgr command:

```
P00>>> wwidmgr -show adapter
wwidmgr available only prior to booting.
Reinit system and try again.
P00>>> init
       .
       .
       .
P00>>> wwidmgr -show adapter
       .
       .
       .
```

For more information on the `wwidmgr` utility, see the *wwidmgr User's Manual*, which is on the Alpha Systems Firmware Update CD-ROM in the `DOC` directory.

4.2.3 Obtain the Fibre Channel Adapter Worldwide Name

A worldwide name (WWN) is a unique number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by the manufacturer prior to shipping. The worldwide name assigned to a subsystem never changes. We recommend that you obtain and record the worldwide names of Fibre Channel components in case you need to verify their target ID mappings in the operating system.

Fibre Channel devices have a 128-bit worldwide ID or name (WWID, WWN) made up of a 64-bit node name (node worldwide name, node WWN, or WWNN) and a 64-bit port name (port worldwide name, port WWN, or WWPN). The most significant four digits are 2000 for the node worldwide name, and 1000 for the port worldwide name for most Fibre Channel adapters currently supported.

There are multiple ways to obtain a Fibre Channel adapter node or port WWN:

- You can obtain the WWN from a label on the Fibre Channel adapter module before you install it.

  Note

  The label on the Fibre Channel adapter module provides the least significant 12 hex digits of both the node and port WWN. Software adds the most significant 4 hex digits (either 2000 for the node WWN or 1000 for the port WWN).

- You can use the `show dev` command to obtain the port WWN (labeled WWN) as follows:

```
P00>>> show dev
       .
       .
       .
pga0.0.0.7.0       PGA0    WWN 1000-0000-c927-242e
pgb0.0.0.8.0       PGB0    WWN 1000-0000-c927-2e67
```

- You can use the `wwidmgr -show adapter` command as follows to obtain the port WWN:

```
P00>>> wwidmgr -show adapter
item       adapter          WWN               Cur. Topo  Next Topo
[ 0]    pga0.0.0.7.0    1000-0000-c927-242e      FABRIC     FABRIC
[ 1]    pgb0.0.0.8.0    1000-0000-c927-2e67      FABRIC     FABRIC
[9999] All of the above.
```

Obtain the Fibre Channel host bus adapter port WWN using the `wwidmgr -show port` command as follows:

```
P00>>> wwidmgr -show port
pga0.0.0.7.0 Link is down.
pgb0.0.0.8.0 Link is down.
[0] 1000-0000-c927-242e
[1] 1000-0000-c927-2e67
```

Note

Use the `wwidmgr -show port` command before connecting the Fibre Channel host bus adapters to the Fibre Channel switches. When executed after the fiber-optic cables are installed, the `wwidmgr -show port` command displays all Fibre Channel host bus adapters connected to the Fibre Channel switch, the Fibre Channel switch ports, and the controller ports.

Record the port WWN of each Fibre Channel adapter for later use in Table A-1. Table 4-1 lists the WWNs for the systems used in the examples in this chapter.

Table 4-1: Fibre Channel Host Bus Adapter Port Worldwide Names

Cluster Member	Fibre Channel Adapter	Port Worldwide Name
Member 1 (`swiss`)	`emx0`	1000-0000-c927-242e
	`emx1`	1000-0000-c927-2e67
Member 2 (`rye`)	`emx0`	1000-0000-c92c-2c81
	`emx1`	1000-0000-c92c-2c10

Note

The examples in this chapter are for a two-node cluster. The MSA1000 will support four-node clusters.

4.3 Install and Set Up the MSA1000 Storage System

If not already completed, install the MSA1000 in the rack, using the documentation that came with the MSA1000 and follow these steps:

1. If necessary, upgrade the MSA1000 firmware to the latest firmware version, which at the time this technical update was re-released, is v4.24.
2. Install additional disk enclosures. You can install two StorageWorks Model 4314R or one StorageWorks Model 4354R disk enclosure.
3. Install disks in the disk enclosures.
4. Install cables between the MSA1000 SCSI I/O module and the Model 4314R or 4354R disk enclosure I/O modules. The cables are 68-pin very high-density cable interconnect (VHDCI-to-VHDCI) cables. Use the following cables:
   - 3X-BC56J-02, 341174-B21 (1.8 m (6 feet))
   - 3X-BC56J-03, 341175-B21 (3.7 meters (12 feet))
   - 3X-BC56J-04, 164604-B21 (7.3 meters (24 feet))

4.4 Installing Fiber-Optic Cables to the MSA1000

The types of fiber-optic cables used will be determined by the type of Fibre Channel host bus adapter and Fibre Channel switch.

The fiber-optic cables used will have either subscriber connector (SC) or Lucent connector (LC) connectors, and the cables may be SC-to-SC, SC-to-LC, or LC-to-LC, depending upon the hardware used. The 1-Gb hardware uses the SC connector while the 2-Gb hardware uses the LC connector.

If the MSA1000 has installed Fibre Channel I/O modules, install the fiber-optic cables as follows (see Figure 3-3):

- For each member system, install an SC-to-SC, SC-to-LC, or LC-to-LC fiber-optic cable between the first Fibre Channel host bus adapter and the first Fibre Channel switch.
- For each member system, install a second SC-to-SC, SC-to-LC, or LC-to-LC fiber-optic cable between the second Fibre Channel host bus adapter and the second Fibre Channel switch.
- Install an SC-to-LC or LC-to-LC fiber-optic cable between the first Fibre Channel switch and the left-most Fibre Channel I/O module (as viewed from the rear of the MSA1000).
- Install a second SC-to-LC or LC-to-LC fiber-optic cable between the second Fibre Channel switch and the right-most Fibre Channel I/O module.

If the MSA1000 has installed MSA SAN Switch 2/8s, install the fiber-optic cables as follows (see Figure 3-4):

- Install an SC-to-LC or LC-to-LC fiber-optic cable between the first Fibre Channel host bus adapter of each member system and the leftmost MSA SAN Switch 2/8.
- Install a second SC-to-LC or LC-to-LC fiber-optic cable between the second Fibre Channel host bus adapter of each member system and the rightmost MSA SAN Switch 2/8.

4.5 Installing Disk Enclosure Power Cables

Install the power cords for the Model 4314R or 4354R disk enclosures. There are no power switches on these disk enclosures; the power is applied to the shelves when you install the power cords.

4.6 Power on the MSA1000 Storage System

Connect the AC power cord for each MSA1000 controller to separate AC power sources. Turn on power to the controller with the power switch located at the right-front of the controller.

The controller performs power-on-self-test (POST) diagnostics and an initialization sequence. It has completed POST and initialization when `MSA1000 Startup Complete` is displayed on the liquid crystal display (LCD) screen on the front of the MSA1000 controller.

4.7 Prepare for CLI Use

Use the command-line interface (CLI) to configure and manage the MSA1000, and to display its status.

To prepare to use the CLI, follow these steps:

1. Install the custom serial cable, part number 259992-001, between a terminal or laptop computer and one of the MSA1000 controllers. This cable has an RJ-45Z (special 10-pin RJ45) connector on one end and a DB-9 connector on the other end.
2. Set up a terminal, of if you are using a laptop, set up a terminal emulator, for example HyperTerminal, as follows:
   - Speed: 19.2 Kb/sec
   - Data bits: 8
   - Parity: None
   - Stop Bits: 1

- Flow Control: None

After the terminal, or terminal emulator is set up, type an Enter or two to obtain the `CLI>` prompt.

The MSA1000 accepts CLI commands in upper- or lowercase. All commands will be typed in uppercase in this document.

See the *hp StorageWorks modular SAN array 1000 user guide* for more information on CLI commands than is provided in this technical update.

Note

MSA1000 commands must be typed completely; they cannot be abbreviated.

4.8 Preparing the Storage for Tru64 UNIX and TruCluster Server Software Installation

Storage for the example two-member TruCluster Server cluster configuration is allocated as shown in Table 4-2.

Table 4-2: Storage Considerations for Example Installation

File System or Disk	Size	Unit Number	Unit ID	Device Name[a]	dsk*n*[b]
Tru64 UNIX disk	2 GB	1	101	`dga101.1001.0.7.0`	73
Cluster root (`/`)	2 GB	2	102	N/A[c]	74
Cluster `/usr`	8 GB	3	103	N/A[c]	75
Cluster `/var`	6 GB	4	104	N/A[c]	76
Quorum disk	1 MB	5	105	N/A[c]	77
Member 1 boot disk	3 GB	6	106	`dga106.1001.0.7.0`	78
Member 2 boot disk	3 GB	7	107	`dga107.1001.0.7.0`	79

[a] The device names are defined in Section 4.9.
[b] The dsk*n* names must be added to the table when they are discovered during the installation (see Section 4.10.1), or by use of the `hwmgr` command as shown in Section 4.10.3.
[c] These units are not assigned an alias for the device unit number by the WWID manager command; therefore, they do not get a device name and will not show up in a console `show dev` display.

As you proceed through the installation, fill out Table A-2 with information specific for your cluster.

To set up storage for a Tru64 UNIX and TruCluster Server configuration using storage as shown in Table 4-2, follow these steps:

1. Change the MSA1000 prompt if you want to. You can use up to 25 alphanumeric characters. The `SET PROMPT` command shown in Example 4-3 changes the MSA1000 prompt to `MSA1000-1>` .

Example 4-3: Changing the MSA1000 Controller Prompt

```
CLI> SET PROMPT "MSA1000-1"
MSA1000-1>
```

Note

Do not add the right angle bracket (⟩) and space to the command; they are added by the `set prompt` command.

2. The MSA1000 creates a connection for each Fibre Channel host bus adapter connected. Use the `SHOW CONNECTIONS` CLI command to verify each of the connections as shown in Example 4-4. Find a connection with a worldwide node name (WWNN) and worldwide port name (WWPN) that matches each of the Fibre Channel host bus adapters in the configuration with the MSA1000. Record the WWPNs in Table A-1.

Example 4-4: Displaying the MSA1000 Connections

```
MSA1000-1> SHOW CONNECTIONS
Connection Name: ⟨Unknown⟩ 1
   Host WWNN = 20000000-C927242E 2
   Host WWPN = 10000000-C927242E 3
   Profile Name = Default 4
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨Unknown⟩
   Host WWNN = 20000000-C9272E67
   Host WWPN = 10000000-C9272E67
   Profile Name = Default
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨Unknown⟩
   Host WWNN = 20000000-C92C2C81
```

Example 4-4: Displaying the MSA1000 Connections (cont.)

```
   Host WWPN = 10000000-C92C2C81
   Profile Name = Default
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨Unknown⟩
   Host WWNN = 20000000-C92C2C10
   Host WWPN = 10000000-C92C2C10
   Profile Name = Default
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online
```

1 The MSA1000 detects all the Fibre Channel host bus adapters that have access to the MSA1000 and creates connections for them. As the connection has not been assigned a name, the name is `unknown`.

There is one information set for each Fibre Channel host bus adapter.

2 The node WWN for the Fibre Channel host bus adapter.

3 The port WWN for the Fibre Channel host bus adapter. For those Fibre Channel host bus adapters that have only one channel, it is the same as the node WWN except for the most significant digit (2 for the node WWN, and 1 for the port WWN).

4 The profile that is set for this connection. The `default` profile is for the `Windows` operating systems and must be changed.

3. A profile name of `default` indicates that the profile for that connection is `Windows`. The profile must be changed to `Tru64`.

   It is difficult to try and keep track of the Fibre Channel node or port WWNs, so you may also want to change the name of the connection from `unknown` to something more meaningful. In Example 4-5, the name of each connection is changed, and the profile is changed to `Tru64`. Obtain the port WWN for each Fibre Channel host bus adapter from Table A-1. The names used in Example 4-5 are taken from Table 4-1.

Example 4-5: Changing the Connection Profile to Tru64

```
MSA1000-1> ADD CONNECTION swiss-emx0 WWPN=10000000-C927242E PROFILE=Tru64
MSA1000-1> ADD CONNECTION swiss-emx1 WWPN=10000000-C9272E67 PROFILE=Tru64
MSA1000-1> ADD CONNECTION rye-emx0 WWPN=10000000-C92C2C81 PROFILE=Tru64
MSA1000-1> ADD CONNECTION rye-emx1 WWPN=10000000-C92C2C10 PROFILE=Tru64
```

Use the SHOW CONNECTIONS command once again to verify that the connection name and profile has been changed, as shown in Example 4-6.

Example 4-6: Verifying the Change in Connection Name and Profile

```
MSA1000-1> SHOW CONNECTIONS
Connection Name: ⟨swiss-emx0⟩
   Host WWNN = 20000000-C927242E
   Host WWPN = 10000000-C927242E
   Profile Name = Tru64
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨swiss-emx1⟩
   Host WWNN = 20000000-C9272E67
   Host WWPN = 10000000-C9272E67
   Profile Name = Tru64
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨rye-emx0⟩
   Host WWNN = 20000000-C92C2C81
   Host WWPN = 10000000-C92C2C81
   Profile Name = Tru64
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online

Connection Name: ⟨rye-emx1⟩
   Host WWNN = 20000000-C92C2C10
   Host WWPN = 10000000-C92C2C10
   Profile Name = Tru64
   Unit Offset = 0
   Controller 2 Port 1 Status = Online
   Controller 1 Port 1 Status = Online
```

4. Determine if any storage units (LUNs) are defined on this MSA1000. There may be other systems using this MSA1000, or you may have

inherited an MSA1000 that was used previously. Use the SHOW UNIT command, as shown in Example 4-7, to determine if there are any existing LUNs.

Example 4-7: Using the SHOW UNIT Command to Display Existing LUNs

```
MSA1000-1> SHOW UNIT
.
.
.
```

If there are existing LUNs that you want to remove, delete one LUN at a time with the DELETE UNIT command, as shown in Example 4-8.

Note

If more than one LUN has been created, only the last LUN created may be deleted. The MSA1000 adds LUNs in numerical order, so only the LUN with the highest number can be deleted.

If you try to delete a unit other than the last unit created, you will be prevented from doing so.

If you try to delete multiple units, for example units 1-7, you can not do it even if storage unit 7 is the highest number storage unit.

If you need to start with a clean MSA1000 with no units created, you can delete all units with the DELETE UNIT ALL -Y command.

Example 4-8: Deleting a LUN

```
MSA1000-1> DELETE UNIT 7
Data will be lost after the unit is deleted.
Do you still want to DELETE unit 7? (Y/N)? Y
Please wait while unit 7 is being deleted...
Unit 7 is deleted successfully.
```

5. Display a list of all installed disks recognized by the MSA1000, as shown in Example 4-9.

Example 4-9: Displaying the Installed Disks

```
MSA1000-1> SHOW DISKS
Disk List: (box,bay) (bus,ID)     Size     Units
 Disk101     (1,01)    (0,00)     9.1GB    none
 Disk102     (1,02)    (0,01)     9.1GB    none
 Disk103     (1,03)    (0,02)     9.1GB    none
 Disk104     (1,04)    (0,03)     9.1GB    none
 Disk105     (1,05)    (0,04)     9.1GB    none
 Disk106     (1,06)    (0,05)     9.1GB    none
 Disk107     (1,07)    (0,08)     9.1GB    none
 Disk108     (1,08)    (1,00)     9.1GB    none
 Disk109     (1,09)    (1,01)     9.1GB    none
 Disk110     (1,10)    (1,02)     9.1GB    none
 Disk111     (1,11)    (1,03)     9.1GB    none
 Disk112     (1,12)    (1,04)     9.1GB    none
 Disk113     (1,13)    (1,05)     9.1GB    none
 Disk114     (1,14)    (1,08)     9.1GB    none
```

6. Create the storage units that will be used for the Tru64 UNIX and TruCluster Server filesystems. Example 4-10 uses the information from Table 4-2 and four 9.2-GB disks to provide RAID0+1. If you select RAID1 (mirroring) and have more than one mirrored pair (4, 6, 8, and so forth disks) data is automatically striped to provide RAID0+1.

Example 4-10: Creating the Storage Units (LUNs)

```
MSA1000-1> ADD UNIT 1 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=2GB [1]
First volume to be configured on these drives.
The logical unit size has been adjusted by  1MB for optimal performance.
Logical Unit size                =    2047 MB
RAID overhead                    =    2047 MB
Total space occupied by new unit  =    4094 MB
Free space left on this volume:   =  134827 MB [2]
Unit 1 is created successfully.

MSA1000-1> ADD UNIT 2 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=2GB [3]
The logical unit size has been adjusted by  1MB for optimal performance.
Logical Unit size                =    2047 MB
RAID overhead                    =    2047 MB
Total space occupied by new unit  =    4094 MB
Free space left on this volume:   =  130732 MB
Unit 2 is created successfully.

MSA1000-1> ADD UNIT 3 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=8GB [4]
The logical unit size has been adjusted by  3MB for optimal performance.
Logical Unit size                =    8189 MB
RAID overhead                    =    8189 MB
Total space occupied by new unit  =   16378 MB
Free space left on this volume:   =  114354 MB
Unit 3 is created successfully.

MSA1000-1> ADD UNIT 4 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=6GB [5]
The logical unit size has been adjusted by  2MB for optimal performance.
Logical Unit size                =    6142 MB
RAID overhead                    =    6142 MB
Total space occupied by new unit  =   12284 MB
Free space left on this volume:   =  102068 MB
Unit 4 is created successfully.
```

Example 4-10: Creating the Storage Units (LUNs) (cont.)

```
MSA1000-1> ADD UNIT 5 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=1MB [6]
A unit cannot be smaller than 15MB.
.
.

MSA1000-1> ADD UNIT 5 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=15MB [6]
MSA1000-1> ADD UNIT 6 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=3GB [7]
.
.

MSA1000-1> ADD UNIT 7 RAID_LEVEL=1 DATA="DISK101 DISK102 DISK108 DISK109" SIZE=3GB [8]
.
.
```

[1] Create the storage unit for the Tru64 UNIX disk. Disks are identified by box and bay numbers. See Figure 1-1. Disk101 is in box 1 (part of the MSA1000) in bay 1. Disk108 is in bay 8 of box 1. Disk numbers must be enclosed in double quotes. You may select a range of sequential disks, for example `DISK101-DISK104`. You can designate the LUN size by megabytes or gigabytes. You can also designate stripe size (for example, STRIPE_SIZE=64) if the default is not sufficient. The default stripe size for RAID0 and 1 is 128 KB. The default stripe size for RAID5 and ADG is 16 KB.

The cache is enabled by default; you can disable it with `CACHE=DISABLE`. Other choices for the `RAID_LEVEL` option are `0`, `5`, and advanced data guarding (`ADG`).

[2] Displays the amount of free space left on the volume.

Note

If you ever need to determine the amount of free space left on a volume, and did not record it when you added the last storage unit, you have to add another storage unit to the volume. There is no command to show volume free space.

[3] Create the storage unit for the cluster root (`/`) file system.

[4] Create the storage unit for the cluster `/usr` file system.

[5] Create the storage unit for the cluster `/var` file system.

[6] Create the storage unit for the quorum disk. An attempt is made to create a storage unit of 1 MB. The MSA1000 will not allow a unit smaller than 15 MB.

7 Create the boot disk for member 1.

8 Create the boot disk for member 2.

Note

You should fill in Table A-2 as you progress.

7. If you want to, add a spare disk to protect against a failed disk. If you add a spare disk to storage unit 1, that spare is assigned to any other storage units that use the same disks as storage unit 1, in our case, disks 101, 102, 108, and 109. You can assign one disk drive to be a spare for several arrays of disks. Add a spare disk, as shown in Example 4-11.

Example 4-11: Adding a Spare Disk

```
MSA1000-1> ADD SPARE UNIT=1 DISK103
Spare drive(s) has been added. Use 'show unit 1' to confirm.
Spare drive(s) has been added. Use 'show unit 2' to confirm.
Spare drive(s) has been added. Use 'show unit 3' to confirm.
Spare drive(s) has been added. Use 'show unit 4' to confirm.
Spare drive(s) has been added. Use 'show unit 5' to confirm.
Spare drive(s) has been added. Use 'show unit 6' to confirm.
Spare drive(s) has been added. Use 'show unit 7' to confirm.
```

8. Enter the SHOW UNIT CLI command to show all units or SHOW UNIT nn (where nn is the unit number) to verify that the information for each storage unit is correct, as shown in Example 4-12.

Example 4-12: Displaying Information About the Storage Unit

```
MSA1000-1> SHOW UNIT 1
Unit 1:
In PDLA mode, Unit 1 is Lun 2; In VSA mode, Unit 1 is Lun 1. [1]
Unit Identifier   : 2 [2]
Device Identifier : 600805F3-000127E0-00000000-CA7B0075 [3]
Cache Status      : Enabled [4]
Max Boot Partition: Enabled [5]
Volume Status     : VOLUME OK [6]
Mirror Init Status: waiting for first write [7]
4 Data Disk(s) used by lun 1:  [8]
   Disk101: Box 1, Bay 01, (SCSI bus 0, SCSI id  0)
   Disk102: Box 1, Bay 02, (SCSI bus 0, SCSI id  1)
   Disk108: Box 1, Bay 08, (SCSI bus 1, SCSI id  0)
   Disk109: Box 1, Bay 09, (SCSI bus 1, SCSI id  1)
Spare Disk(s) used by lun 1:
```

Example 4-12: Displaying Information About the Storage Unit (cont.)

```
   Disk103: Box 1, Bay 03, (SCSI bus 0, SCSI id  2) [9]
Logical Volume Raid Level: MIRROR FAULT TOLERANCE (Raid 1) [10]
                           stripe_size=128kB [11]
Logical Volume Capacity : 2,047MB [12]
MSA1000-1> SHOW UNIT 7
Unit 7:
In PDLA mode, Unit 7 is Lun 8; In VSA mode, Unit 7 is Lun 7.
Unit Identifier   : 8
Device Identifier : 600805F3-000127E0-00000000-ACB70082
Cache Status      : Enabled
Max Boot Partition: Enabled
Volume Status     : VOLUME OK
Mirror Init Status: waiting for first write
4 Data Disk(s) used by lun 7:
   Disk101: Box 1, Bay 01, (SCSI bus 0, SCSI id  0)
   Disk102: Box 1, Bay 02, (SCSI bus 0, SCSI id  1)
   Disk108: Box 1, Bay 08, (SCSI bus 1, SCSI id  0)
   Disk109: Box 1, Bay 09, (SCSI bus 1, SCSI id  1)
Spare Disk(s) used by lun 7:
   Disk103: Box 1, Bay 03, (SCSI bus 0, SCSI id  2)
Logical Volume Raid Level: MIRROR FAULT TOLERANCE (Raid 1)
                           stripe_size=128kB
Logical Volume Capacity : 3,067MB
```

[1] Peripheral device LUN addressing (PDLA) mode is used by all operating systems supporting the MSA1000 except HP-UX. In PDLA mode, the units are numbered starting with 1, but the first LUN is 2.

Volume set addressing (VSA) mode is used by HP-UX. The unit and LUN are numbered the same.

For Tru64 UNIX, PDLA mode is used.

[2] The Unit Identifier (Unit ID) is assigned the same number as the LUN.

[3] The Device Identifier for the logical volume. It is equivalent to the node WWN and port WWN. It is assigned by the MSA1000.

[4] The MSA1000 is enabled to use cache, which is the default.

[5] `Max Boot Partition` is ignored by all operating systems except Windows.

[6] Describes the current condition of the volume.

[7] Provides the parity or mirror status, depending on the RAID level selected. Parity calculation does not start until the first write is

performed on the drive. Mirror status is not provided until after the first write.

8 Lists the disks used by this LUN.

9 A list of spare disks assigned to this LUN.

10 The RAID level.

11 The stripe size.

12 The total capacity of this logical volume

9. When you create the unit with the `ADD UNIT` command, the unit ID is automatically assigned to the LUN number as can be seen in Example 4-12.

   If you have, or plan on having, multiple MSA1000 storage systems being accessed by your TruCluster Server member systems, set the unit ID to ensure that you do not have any problems identifying them from the Tru64 UNIX operating system.

Note

The unit ID is used by the `wwidmgr` utility to make the device available to the console as a boot device, and shows up as the `IDENTIFIER` during a Tru64 UNIX installation, or when using the `hwmgr view device` command.

Caution

The *hp StorageWorks modular SAN array 1000 user guide* provides examples of setting the unit ID to a sequence of alphabetical characters. If you set the unit ID to any non-numerical value, the `wwidmgr` utility will not recognize the storage unit, and it will not be displayed as a valid device for Tru64 UNIX software installation.

You can create a maximum of 32 LUNs, so you could change the unit IDs for an MSA1000 in groups of 32 numbers. It will be easier to recognize the LUNs if you change the unit IDs on those LUNs on the first MSA1000 to numbers in the range 100 to 199. The LUNs on the second MSA1000 could be assigned unit IDs in the range 200 to 299, and so on.

The unit IDs of the storage units created in Example 4-10 will be changed to the unit IDs listed in Table 4-2 by the commands shown in Example 4-13.

Example 4-13: Changing a LUN Unit ID

```
MSA1000-1> SET UNIT_ID 1 101 [1]
MSA1000-1> SET UNIT_ID 2 102 [2]
MSA1000-1> SET UNIT_ID 3 103 [3]
MSA1000-1> SET UNIT_ID 4 104 [4]
MSA1000-1> SET UNIT_ID 5 105 [5]
MSA1000-1> SET UNIT_ID 6 106 [6]
MSA1000-1> SET UNIT_ID 7 107 [7]
```

[1] Change the unit ID of storage unit 1, LUN 2 to unit ID 101.

[2] Change the unit ID of storage unit 2, LUN 3 to unit ID 102.

[3] Change the unit ID of storage unit 3, LUN 4 to unit ID 103.

[4] Change the unit ID of storage unit 4, LUN 5 to unit ID 104.

[5] Change the unit ID of storage unit 5, LUN 6 to unit ID 105.

[6] Change the unit ID of storage unit 6, LUN 7 to unit ID 106.

[7] Change the unit ID of storage unit 7, LUN 8 to unit ID 107.

10. If you have any other TruCluster Server cluster, or any other system accessing storage units on this MSA1000 storage system, you must restrict access to cluster storage to the cluster member systems. This is generally called selective storage presentation (SSP). The CLI uses an access control list (ACL) to control access to storage units.

    The `ADD ACL` command provides access to the named storage unit based on the connection name or port WWN. You can add storage units to the ACL using single unit numbers (`UNIT=1`), a range of unit numbers (`UNIT=1-7`), or all units (`UNIT=ALL`).

    Example 4-14 provides an example of adding all four of the TruCluster Server Fibre Channel host bus adapters to the access control list for each of the storage units created in Example 4-10.

    Note

    The `ADD ACL` command uses the unit numbers, not the LUN number or unit ID.

Example 4-14: Creating an Access Control List for TruCluster Server Storage Units

```
MSA1000-1> ADD ACL CONNECTION=swiss-emx0 UNIT=1-7  [1]
Enabling ACL. Allowing 10000000-C927242E access to unit 1 2 3 4 5 6 7.
MSA1000-1> ADD ACL CONNECTION=swiss-emx1 UNIT=1-7  [2]
Enabling ACL. Allowing 10000000-C9272E67 access to unit 1 2 3 4 5 6 7.
MSA1000-1> ADD ACL WWPN=10000000-C92C2C81 UNIT=1-7 [3]
Enabling ACL. Allowing 10000000-C92C2C81 access to unit 1 2 3 4 5 6 7.
MSA1000-1> ADD ACL WWPN=10000000-C92C2C10 UNIT=1-7 [4]
Enabling ACL. Allowing 10000000-C92C2C10 access
to unit 1 2 3 4 5 6 7.
```

[1] Use the connection name to provide access to storage units 1-7 for Fibre Channel host bus adapter `emx0` on member `swiss`.

[2] Use the connection name to provide access to storage units 1-7 for Fibre Channel host bus adapter `emx1` on member `swiss`.

[3] Use the port WWN to provide access to storage units 1-7 for Fibre Channel host bus adapter `emx0` on member `rye`.

[4] Use the port WWN to provide access to storage units 1-7 for Fibre Channel host bus adapter `emx1` on member `rye`.

11. Use the `SHOW ACL` CLI command as shown in Example 4-15 to verify that the ACL provides access to only the TruCluster Server member systems for the cluster storage units.

Example 4-15: Verifying the Access Control List

```
MSA1000-1> SHOW ACL
    ACL is enabled:
    Connection        WWPN               Units
    swiss-emx0        10000000-C927242E  1, 2, 3, 4, 5, 6, 7
    swiss-emx1        10000000-C9272E67  1, 2, 3, 4, 5, 6, 7
    rye-emx0          10000000-C92C2C81  1, 2, 3, 4, 5, 6, 7
    rye-emx1          10000000-C92C2C10  1, 2, 3, 4, 5, 6, 7
    Inaccessible Units:  None
```

4.9 Setting the Device Unit Number and Verifying Valid Boot Devices

This section covers the remaining steps you must complete before you install the Tru64 UNIX and TruCluster Server software:

1. Set the device unit number of the disk where you will install the base operating system software, and set the device unit number of the first cluster member boot disk. Setting the device unit number allows the console and installation scripts to recognize the disks (see Section 4.9.1).
2. Verify that the console recognizes these disks as valid boot devices (see Section 4.9.2).

4.9.1 Set the Device Unit Number

The device unit number is a subset of the device name as shown in a `show device` console display. For example, in the device name `dga101.1001.0.7.0`, the device unit number is `101` (as in `dga101`). The console uses this device unit number to identify a storage unit. When you set a device unit number, you are really setting an alias for the device worldwide name (WWN, also called the worldwide ID, WWID). The 128-bit WWN is too large to be used as the device unit number, so an alias is used instead. The 128-bit WWN is labeled `Device Identifier` in an MSA1000 `SHOW UNIT` display.

This section describes how to use the `wwidmgr -quickset` command to set the device unit number for the Fibre Channel disks to be used as the Tru64 UNIX installation disk or cluster member system boot disks.

To set the device unit number for a Fibre Channel device, follow these steps:

1. From Table A-2, obtain the unit ID for the virtual disk to be used as the Tru64 UNIX installation disk or cluster member system boot disks. The example installation uses the information in Table 4-2. The unit ID is referred to as the user-defined identifier (UDID) for the console software and WWID manager (`wwidmgr`).

   For example, in Table 4-2 the Tru64 UNIX disk has a unit ID of 101. The unit ID for the cluster member 1 boot disk is 106, and the cluster member 2 boot disk is 107.

2. From the AlphaServer console, use the `wwidmgr -clear all` command to clear the stored Fibre Channel `wwid1`, `wwid2`, `wwid3`, `wwid4`, `N1`, `N2`, `N3`, and `N4` console environment variables. You want to start with all `wwid`*n* and `N`*n* variables clear.

   A console initialization is generally required before you can use the `wwidmgr` command. For example:

```
P00>>> init
    .
    .
    .
P00>>> wwidmgr -clear all
P00>>> show wwid*
wwid0
wwid1
wwid2
wwid3

P00>>> show n*
N1
N2
N3
N4
```

Note

The console only creates devices for which the `wwid`*n* console environment variable has been set, and that are accessible through an MSA1000 N_Port as specified by the `N`*n* console environment variable also being set. These console environment variables are set with the `wwidmgr -quickset` or `wwidmgr -set wwid` commands. The use of the `wwidmgr -quickset` command is shown in the next step.

3. Use the `wwidmgr` command with the `-quickset` option to set a device unit number for the Tru64 UNIX installation disk and the first cluster member system boot disk.

   The `wwidmgr` command with the `-quickset` option is used to define a device unit number, based on the unit ID, as an alias for the WWID for the Tru64 UNIX installation disk and the first cluster member system boot disk. The `wwidmgr -quickset` utility sets the device unit number and also provides a display of the device names and how the disk is reachable (reachability display).

   The `wwidmgr -quickset` command may generate multiple device names for a given device unit number, because each possible path to a storage unit is given its own device name.

   Set the device unit number for the Tru64 UNIX installation disk and the first cluster member system boot disk as follows:

   a. Set the device unit number for the Tru64 UNIX installation disk to 101 (the same as the unit ID), as shown in Example 4-16.

Example 4-16: Setting the Device Unit Number for the BOS Installation Disk

```
P00>>> wwidmgr -quickset -udid 101

Disk assignment and reachability after next initialization:

6008-05f3-0001-27e0-0000-0000-ca7b-0075 [1]
                      via adapter:          via fc nport:      connected: [2]
dga101.1001.0.7.0 [3]   pga0.0.0.7.0 [4]    5008-05f3-0001-bd51 [5]   Yes
dga101.1002.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd59       No

dgb101.1001.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd51       Yes
dgb101.1002.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd59       No
```

The `wwidmgr -quickset` command provides a reachability display equivalent to entering the `wwidmgr -show reachability` command. The reachability part of the display provides the following information:

[1] The worldwide name (WWN) for the storage unit that is to be accessed. This WWN is assigned to a `wwid`*n* variable.

[2] In the `connected` column, whether the storage unit is currently available through the Fibre Channel host bus adapter to controller port connection.

[3] The new device name for the storage unit.

[4] The Fibre Channel adapters through which a connection to the storage unit is potentially available.

[5] The port WWN of the controller ports (N_Ports) that will be used to access the storage unit. This will show up as an `N`*n* variable.

b. Set the device unit number for the first cluster member system boot disk to 106, as shown in Example 4-17.

Example 4-17: Setting the Device Unit Number for the First Cluster Member Boot Disk

```
P00>>> wwidmgr -quickset -udid 106

Disk assignment and reachability after next initialization:

6008-05f3-0001-27e0-0000-0000-ca7b-0075
                      via adapter:          via fc nport:      connected:
dga101.1001.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd51      Yes
dga101.1002.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd59      No

dgb101.1001.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd51      Yes
dgb101.1002.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd59      No
```

Example 4-17: Setting the Device Unit Number for the First Cluster Member Boot Disk (cont.)

```
6008-05F3-0001-27E0-0000-0000-B98B-007A
                          via adapter:         via fc nport:        connected:
dga106.1001.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd51      Yes
dga106.1002.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd59       No
dgb106.1001.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd51       Yes
dgb106.1002.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd59       No
```

Note

After the second member system has been added to the cluster, you have to set the device unit number on that system before booting.

c. A console initialization is required to exit the `wwidmgr`, and to make the device names available to the console `show dev` command:

```
P00>>> init
     .
     .
     .
```

The device names have now been set for the Tru64 UNIX disk and first cluster member system boot disks.

In the reachability portion of the display, each storageset is reachable, although not necessarily connected, from Fibre Channel adapter `pga` through two controller ports and from Fibre Channel adapter `pgb` through two controller ports. Also, the device unit number has been set for each Fibre Channel host bus adapter to controller port connection, even if the storage unit is not currently reachable (connected) through that connection.

4.9.2 Displaying Valid Boot Devices

The only Fibre Channel devices that are displayed by the console `show dev` command are those devices that have been assigned to a `wwid`*n* environment variable with the `wwidmgr -quickset` command.

Any device shown in the reachability display can be used as a boot device. The `bootdef_dev` console environment variable can be set to any, or several, of these devices. Also, the cluster installation script sets the `bootdef_dev` console environment variable to up to four of these devices.

If you enter the `show wwid*` console command now, it will show that the environment variable `wwid`*n* is set for two logical volumes. Also, the `show`

n* command shows that the units are accessible through two controller N_Ports, as shown in Example 4-18.

Example 4-18: Displaying the wwid and n Console Environment Variables

```
P00>>> show wwid*
wwid0                101 1 WWID:01000010:6008-05f3-0001-27e0-0000-0000-ca7b-0075
wwid1                106 1 WWID:01000010:6008-05F3-0001-27E0-0000-0000-B98B-007A
wwid2
wwid3
P00>>>show n*
N1                   500805f30001bd51
N2                   500805f30001bd59
N3
N4
```

Example 4-19 provides sample device names as displayed by the show dev command after using the wwidmgr -quickset command to set the device unit numbers. These devices are available to use as boot devices.

Example 4-19: Sample Fibre Channel Device Names

```
 P00>>> show dev
dga101.1001.0.7.0         $1$DGA101   COMPAQ MSA1000 VOLUME  3.02
dga101.1002.0.7.0         $1$DGB101   COMPAQ MSA1000 VOLUME  3.02
dgb101.1001.0.8.0         $1$DGA101   COMPAQ MSA1000 VOLUME  3.02
dgb101.1002.0.8.0         $1$DGB101   COMPAQ MSA1000 VOLUME  3.02
dga106.1001.0.7.0         $1$DGA106   COMPAQ MSA1000 VOLUME  3.02
dga106.1002.0.7.0         $1$DGB106   COMPAQ MSA1000 VOLUME  3.02
dgb106.1001.0.8.0         $1$DGA106   COMPAQ MSA1000 VOLUME  3.02
dgb106.1002.0.8.0         $1$DGB106   COMPAQ MSA1000 VOLUME  3.02
dka500.5.0.2000.0          DKA500                     RRD47  1206
dkb0.0.0.2001.1            DKB0                    RZ1CD-CS  0306
.
.
.

pga0.0.0.7.0            PGA0        WWN 1000-0000-c927-242e
pgb0.0.0.8.0            PGB0        WWN 1000-0000-c927-2e67
.
.
.
```

Note

The only Fibre Channel devices displayed by the console show dev command are those devices that have been assigned to a wwid*n* environment variable.

At this point you are ready to install the Tru64 UNIX operating system and TruCluster Server software.

4.10 Installing the Software

This section covers the installation of the Tru64 UNIX and TruCluster Server software:

1. Install the base operating system (see Section 4.10.1).
2. If you are not installing the TruCluster Server software, set the `bootdef_dev` console environment variable to ensure that there is a path to the boot disk if the RAID array controllers have failed over (see Section 4.10.2). If you are installing TruCluster Server software, the `bootdef_dev` console environment variable is set by the installation script.
3. Determine the `dsk`*n* to use for cluster installation (see Section 4.10.3).
4. Label the disks to be used for cluster installation (see Section 4.10.4).
5. Install the TruCluster Server software (see Section 4.10.5).
6. Add additional cluster members (see Section 4.10.6).

4.10.1 Install the Base Operating System

Using the Tru64 UNIX *Installation Guide* as a reference, boot from the CD-ROM and perform a full installation of the Tru64 UNIX Version 5.1A or 5.1B operating system.

When the installation procedure displays the list of disks that are available for operating system installation as shown in Example 4-20, look for the identifier in the Location column. Verify the identifier from Table A-2. For the example being followed in this document, use Table 4-2.

Example 4-20: Disks Available for Software Installation

```
Select a disk for the root file system.  The
root file system will be placed on the "a" partition of the disk
you choose.

To visually locate a disk, enter "ping <disk>",
where <disk> is the device name (for example, dsk0) of the disk you
want to locate.  If that disk has a visible indicator light, it will
blink until you are ready to continue.

      Device      Size  Controller  Disk
      Name       in GB  Type        Model          Location
1)    dsk0         4.0  SCSI        RZ2CA-LA       bus-0-targ-0-lun-0
2)    dsk1         4.0  SCSI        RZ2CA-LA       bus-0-targ-1-lun-0
3)    dsk2         4.0  SCSI        RZ1CB-CS       bus-1-targ-0-lun-0
4)    dsk3         4.0  SCSI        RZ1CF-CF       bus-1-targ-1-lun-0
5)    dsk4         1.0  SCSI        RZ26           bus-1-targ-2-lun-0
6)    dsk5         2.0  SCSI        RZ28           bus-1-targ-4-lun-0
7)    dsk6         1.0  SCSI        RZ26L          bus-1-targ-5-lun-0
8)    dsk73        2.0  SCSI        MSA1000        IDENTIFIER=101
```

Example 4-20: Disks Available for Software Installation (cont.)

```
9)    dsk78       3.0  SCSI       MSA1000        IDENTIFIER=106
```

Record the `/dev/disk/dsk`*n* value for the Tru64 UNIX disk in Table A-2. For the example we are using, identifier 101 is `dsk73` (see Example 4-20 and Table 4-2.)

Complete the installation, following the instructions in the Tru64 UNIX *Installation Guide* for your version of the operating system.

If you are only installing the base operating system, and not installing TruCluster Server, set the `bootdef_dev` console environment variable to multiple paths before you boot the operating system (see Section 4.10.2).

Note

Ensure that you install the patch kit needed for your version of the operating system.

If you are installing TruCluster Server software, skip Section 4.10.2 and proceed directly to Section 4.10.3.

4.10.2 Reset the bootdef_dev Console Environment Variable

After installing the operating system software, shut down the operating system to the console prompt. Use the console `show device` command to verify that the `bootdef_dev` console environment variable is set to select multiple paths to the boot device and not just one path.

If it is set to select only one path to the boot device, set it to select multiple paths, as follows:

1. Examine the reachability display provided by the `wwidmgr -show reachability` command for the device names that can access the storage unit from which you are booting.

Notes

Choose device names that show up as both `Yes` and `No` in the reachability display `connected` column. Note, that for the MSA1000, only one controller is normally active for a storage unit. You must ensure that the unit is reachable if the controllers have failed over.

If you have an NSPOF configuration, use device names for at least two host bus adapters.

For example, to ensure that you have a connected boot path in case of a failed host bus adapter or controller failover, choose device names for multiple host bus adapters and each controller port. If you use the reachability display for member system 1's boot disk as shown in Example 4-17, choose all of the following device names when setting the `bootdef_dev` console environment variable for the first cluster member system:

```
dga101.1001.0.7.0
dga101.1002.0.7.0
dgb101.1001.0.8.0
dgb101.1002.0.8.0
```

If the `bootdef_dev` console environment variable ends up with all boot paths in an unconnected state, you can use the `ffauto` or `ffnext` console environment variables to force a boot device from a `not connected` to a `connected` state.

The `ffauto` console environment variable is effective only during autoboots (boots other than manual boots). Use the `set ffauto on` console command to enable `ffauto`. (The default for `ffauto` is `off`.) It is stored in nonvolatile memory so it persists across system resets and power cycles.

During an autoboot, the console attempts to boot from each connected device listed in the `bootdef_dev` console environment variable. If `ffauto` is `on`, and if the end of devices listed in `bootdef_dev` is reached without successfully booting, the console starts again at the beginning of the devices listed in the `bootdef_def` console environment variable. This time, devices that are not connected are changed to `connected` and an attempt is made to boot from that device.

The `ffnext` console environment variable is a one-time variable. It does not persist across a system reset, power cycle, or reboot. This variable may be used (`set ffnext on`) to cause the next command to a `not connected` device to change the state to `connected`. After the command has been executed, the `ffnext` variable is automatically set to `off`, so it has no further effect.

For more information on using the `ffauto` and `ffnext` console environment variables, see the *wwidmgr User's Manual*.

2. Set the `bootdef_dev` console environment variable for the base operating system boot disk to a comma-separated list of several of the boot paths that show up in the reachability display (`wwidmgr -show reachability`). You must initialize the system to use any of the device names in the `bootdef_dev` variable as follows:

```
P00>>> set bootdef_dev \
dga101.1001.0.7.0, dga101.1002.0.7.0, \
dgb101.1001.0.8.0, dgb101.1002.0.8.0
POO>>> init
```

Note

The console System Reference Manual (SRM) software guarantees that you can set the `bootdef_dev` console environment variable to a minimum of four device names. You may be able to set it to five, but only four are guaranteed.

4.10.3 Determining /dev/disk/dskn to Use for a Cluster Installation

Before installing the TruCluster Server software, you must determine which `/dev/disk/dsk`*n* to use for the various TruCluster Server disks.

To determine the `/dev/disk/dsk`*n* to use for the cluster disks, follow these steps:

1. On the system where you installed the Tru64 UNIX operating system, if you have not already done so, boot the system. With the Tru64 UNIX operating system at single-user or multi-user mode, use the hardware manager utility (`hwmgr`) with the `view devices` option to display all devices on the system. Pipe the command through the `grep` utility to search for any items with the `IDENTIFIER` or `MSA1000` qualifiers, as shown in Example 4-21 or Example 4-22.

Example 4-21: Using the hwmgr to Display Devices with the Identifier Qualifier

```
# hwmgr view dev | grep IDENTIFIER
 HWID: Device Name          Mfg      Model             Location
 ------------------------------------------------------------------------
   85: /dev/disk/dsk11c     DEC      HSG80             IDENTIFIER=6
  147: /dev/disk/dsk73c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=101
  148: /dev/disk/dsk74c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=102
  149: /dev/disk/dsk75c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=103
  150: /dev/disk/dsk76c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=104
  151: /dev/disk/dsk77c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=105
  152: /dev/disk/dsk78c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=106
  153: /dev/disk/dsk79c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=107
```

Example 4-22: Using the hwmgr to Display Devices with the MSA1000 Qualifier

```
# hwmgr view dev | grep MSA1000
 HWID: Device Name          Mfg      Model             Location
 ------------------------------------------------------------------------
  147: /dev/disk/dsk73c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=101
  148: /dev/disk/dsk74c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=102
  149: /dev/disk/dsk75c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=103
  150: /dev/disk/dsk76c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=104
  151: /dev/disk/dsk77c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=105
  152: /dev/disk/dsk78c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=106
  153: /dev/disk/dsk79c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=107
  160: /dev/cport/scp3                MSA1000           bus-5-targ-6-lun-0
```

If you know that you have set the unit ID for a large number of disks, you can also search for the unit ID, as shown in Example 4-23.

Example 4-23: Using the hwmgr to Display Devices with a Specific Unit ID

```
# hwmgr view dev | grep IDENTIFIER=106
 HWID: Device Name          Mfg      Model             Location
 ------------------------------------------------------------------------
  152: /dev/disk/dsk78c     COMPAQ   MSA1000 VOLUME    IDENTIFIER=106
```

2. Search the display for the identifiers for each of the cluster installation disks and record the `/dev/disk/dsk`*n* values in Table A-2. The disks used for the example installation are shown in Table 4-2.

   If you use the `grep` utility to search for a specific unit ID, for example, `hwmgr view dev | grep IDENTIFIER=106`, repeat the command to determine the `/dev/disk/dsk`*n* for each of the remaining cluster disks. Record the information for use when you install the cluster software.

Note

You must label the disks before you install cluster software.

4.10.4 Label the Disks to Be Used to Create the Cluster

Before you run `clu_create` to create the first cluster member or `clu_add_member` to add subsequent cluster members, you must label the disks to be used for cluster software.

On the system where you installed the Tru64 UNIX operating system, if you have not already done so, boot the system. Determine the `/dev/disk/dsk`*n* values to use for cluster installation (see Table A-2 or Table 4-2).

Initialize disk labels for all disks needed to create the cluster. The example uses disks `dsk74` (cluster root (`/`)), `dsk75` (`/usr`), `dsk76` (`/var`), and `dsk77` (`Quorum`). For example:

```
# disklabel -z dsk74
  disklabel: Disk /dev/rdisk/dsk74c is unlabeled
# disklabel -rw dsk74
```

4.10.5 Install the TruCluster Server Software and Create the First Cluster Member

After labeling the disks, use the TruCluster Server *Cluster Installation* procedures for your version of the operating system and install the TruCluster Server software on the first cluster member (the system where you just installed Tru64 UNIX).

After installing the TruCluster Server software subsets, run the `clu_create` command to create the first cluster member using the procedures in the TruCluster Server *Cluster Installation* manual.

4.10.6 Add an Additional System to the Cluster

To add an additional system to the cluster, follow these steps:

Note

You can have up to eight member systems in a cluster, but only four systems can be connected to the MSA1000.

1. On the system where you installed the Tru64 UNIX operating system and TruCluster Server software, boot the system into the cluster as a single-member cluster.
2. Referring to the TruCluster Server *Cluster Installation* manual procedures, use `clu_add_member` to add a cluster member.

   Before you boot the system being added to the cluster, on the newly added cluster member:

   a. Use the `wwidmgr` utility with the `-quickset` option to set the device unit number for the member system boot disk, as shown in Example 4-24. For member system 2 in the example configuration, it is the storage unit with unit ID 107 (see Table 4-2):

Example 4-24: Setting Device Unit Number for Additional Member System

```
P00>>> wwidmgr -quickset -udid 107

Disk assignment and reachability after next initialization:

6008-05F3-0001-27E0-0000-0000-ACB7-0082
                          via adapter:         via fc nport:        connected:

dga107.1001.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd51      Yes
dga107.1002.0.7.0       pga0.0.0.7.0        5008-05f3-0001-bd59      No

dgb107.1001.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd51      Yes
dgb107.1002.0.8.0       pgb0.0.0.8.0        5008-05f3-0001-bd59      No
P00>>> init
```

   b. Set the `bootdef_dev` console environment variable to one reachable path (`Yes` in the connected column of Example 4-24) to the member system boot disk:

      ```
      P00>>> set bootdef_dev dga107.1001.0.7.0
      ```

   c. Boot `genvmunix` on the newly added cluster member system. Each installed subset will be configured and a new kernel will be built and installed.

3. Boot the new cluster member system into the cluster and complete the cluster installation.

Note

The cluster installation procedure sets the `bootdef_dev` console environment variable to multiple paths.

You should verify that it is set to multiple paths (`show bootdef_dev`).

If it is not set to multiple paths, enter the `wwidmgr -show reachability` command to determine the paths to the boot disk. An example of the reachability display is shown in Example 4-17.

Set the `bootdef_dev` console environment variable, as shown in step 2 of Section 4.10.2.

A

Worldwide Names, File System, Unit Number, Unit ID, Device Name, and Disk Name

Table A-1: Fibre Channel Host Bus Adapter Worldwide Names

Cluster Member	Fibre Channel Adapter	Port Worldwide Name
Member 1	emx0	
	emx1	
Member 2	emx0	
	emx1	
Member 3	emx0	
	emx1	
Member 4	emx0	
	emx1	

Table A-2: Storage Allocation for TruCluster Server Installation

File System or Disk	Size	Unit Number	Unit ID	Device Name	dsk*n*
Tru64 UNIX disk					
Cluster root (/)					
Cluster /usr					
Cluster /var					
Quorum disk					
Member 1 boot disk					
Member 2 boot disk					
Member 3 boot disk					
Member 4 boot disk					

Index

Numbers and Special Characters

164604-B21, 2-3, 4-8
259992-001, 4-9
341174-B21, 2-3, 4-8
341175-B21, 2-3, 4-8
3X-BC56J-02, 2-3, 4-8
3X-BC56J-03, 2-3, 4-8
3X-BC56J-04, 2-3, 4-8

A

access control list
(*See* ACL)
ACL
controlling access with, 4-20
setting, 4-20
verifying, 4-21
active/passive mode, 2-3, 3-7
add ACL command, 4-20
add connection command, 4-12
add spare command, 4-17
add unit command, 4-15

B

boot devices
available, 4-26
boot problem
wwidmgr, 4-5
bootdef_dev console environment variable, 4-25, 4-29, 4-33
resetting, 4-28
setting, 4-28, 4-30, 4-33

C

CLI
terminal setup, 4-9
CLI command
add ACL, 4-20
add connection, 4-12
add spare, 4-17
add unit, 4-15
delete unit, 4-14
set prompt, 4-11
set unit_id, 4-19
show ACL, 4-21
show connections, 4-11
show disks, 4-14
show unit, 4-17
clu_create command, 4-32
command line interface
(*See* CLI, CLI command)
connection
changing the name of, 4-12
verifying, 4-13
console environment variable
bootdef_dev, 4-25, 4-27, 4-28, 4-29, 4-30, 4-33
ffauto, 4-29
ffnext, 4-29
N, 4-23
wwid, 4-23, 4-24

D

delete unit command, 4-14
device name, 4-22
device unit number
setting, 4-22, 4-23

disk enclosures
model 4314R, 1-2
model 4354R, 1-2
disklabel command, 4-32
disks
available for software installation, 4-27
maximum number, 1-2
number of, 1-2

E

Ethernet switch
10Base-T Ethernet connection, 4-2

F

fabric mode, 4-3
failure to boot, 4-5
ffauto console environment variable, 4-29
ffnext console environment variable, 4-29
Fibre Channel
configuring hardware, 4-1
displaying available boot devices, 4-25
fabric mode, 4-3
labeling disks, 4-32
NVRAM, 4-3
NVRAM read failed message, 4-4
setting topology, 4-4
switch installation, 4-2
verifying topology, 4-3
Fibre Channel adapter
installing, 4-2
obtaining the worldwide name of, 4-7
setting to run on fabric, 4-3, 4-4
worldwide name, 4-7
Fibre Channel I/O module, 1-2, 4-2
Fibre Channel switch
installation, 4-2
overview, 4-2
file
/var/adm/messages, 4-7

H

hardware configuration
Fibre Channel host bus adapter, 4-3
verifying FCA topology, 4-3
hwmgr command
-view dev, 4-30, 4-31

I

init command, 4-5, 4-30
initializing
after setting bootdef_dev console environment variable, 4-30
after using wwidmgr, 4-5
installation
Fibre Channel adapter, 4-2

L

labeling installation disks, 4-32

M

Model 4314R disk enclosure, 1-2
Model 4354R disk enclosure, 1-2
Modular Smart Array 1000
(*See* MSA1000)
MSA SAN switch 2/8, 1-2
MSA1000
1-Gb and 2-Gb compatible, 1-2
active/passive mode, 2-3, 3-7
adding a spare disk, 4-17
changing the connection name, 4-12
changing the prompt, 4-11
changing the unit ID, 4-19
creating access control list, 4-20

creating storage units, 4-15
deleting a unit, 4-14
disk drive compatibility, 1-2
displaying installed disks, 4-14
displaying the storage unit, 4-17
FC I/O module, 1-2
Fibre Channel compatibility, 1-2
installing disk enclosure power cables, 4-9
installing fiber-optic cables, 4-8
installing SCSI cables, 4-8
MSA Hub 2/3, 1-2
MSA SAN switch 2/8, 1-2
setting the profile, 4-12
software support for, 1-1
supported Fibre Channel adapters, 2-2
supported Fibre Channel connections, 1-2
supported software versions, 2-1
total disk capacity, 1-2
verifying connections, 4-11
verifying the ACL, 4-21
MSA1000 cluster configuration
not supported, 3-7
NSPOF with FC switches, 3-3
NSPOF with MSA SAN switch 2/8, 3-6
single controller, 3-1

N

new hardware delivery
(*See* NHD7)
NHD7, 2-1
non volatile random access memory
(*See* NVRAM)
NVRAM
read failed message, 4-4
set to fabric mode, 4-3

O

optical cable lengths, 2-3

P

patches, 2-1
PDLA, 4-18
peripheral device LUN addressing
(*See* PDLA)
profile
changing, 4-12
default, 4-12

R

reset
(*See* system reset)

S

SCSI bus IDs, 1-3
SCSI bus locations, 1-3
SCSI I/O module, 1-3
set bootdef_dev command, 4-30, 4-33
set ffauto command, 4-29
set ffnext command, 4-29
set unit_id command, 4-19
show ACL command, 4-21
show config command, 4-6
show connections command, 4-11, 4-13
show dev command, 4-26
show disks command, 4-14
show n command, 4-22, 4-26
show unit command, 4-14, 4-17
show wwid command, 4-22, 4-26
StorageWorks Modular Smart Array 1000
(*See* MSA1000)
supported

Fibre Channel adapters, 2-2
software versions, 2-1
system boot problem, 4-5
system reset
after using wwidmgr, 4-5

T

topology, 4-3, 4-5

U

Ultra3 SCSI, 1-3

V

valid boot devices, 4-25
/var/adm/messages file, 4-7
variable
(*See* console environment variable)
volume set addressing
(*See* VSA)
VSA, 4-18

W

WorldWide ID Manager
(*See* wwidmgr, wwidmgr command)
worldwide name
(*See* WWN)
WWID
(*See* WWN)
wwidmgr
failure to boot, 4-5
unavailable, 4-5
wwidmgr command, 4-23
-clear, 4-22, 4-23
-quickset, 4-22, 4-23, 4-24, 4-25, 4-33
-set adapter, 4-5
-show adapter, 4-5
-show port, 4-7
-show reachability, 4-24, 4-28, 4-30, 4-34
WWN
description, 4-6
obtaining, 4-6